

Processing-Time Effects on PQC-Based Over-the-Air Rekeying over CCSDS TC/TM Links

Ha-Gyeong Kim¹[0009–0008–5840–5517] and Hwa-Jeong Seo¹[0000–0003–0069–9061]

Hansung University, Seoul, Korea
kimhaha4420@gmail.com, hwajeong84@gmail.com

Abstract. When post-quantum cryptography (PQC)-based over-the-air rekeying (OTAR) is evaluated over CCSDS Telecommand (TC) and Telemetry (TM) links, omitting endpoint processing time can change not only the final transport outcome but also the protocol execution path shaped by COP-1 retransmission and CLCW return feedback. This study instantiates the P1–P2–P3 exchange of a PQC-based Space Data Link Security (SDLS) key update over a CCSDS TC/TM transport model and performs paired comparisons between a zero-processing baseline, T_0 , and a processing-aware model, T_1 , whose service times are measured at seven logical stages of a study-specific Triple-KEM reference implementation. The two models are compared at three observation levels—final transport outcome, canonical event path, and continuous timing—while varying processing scale, finite contact time, direction-specific loss, and schedule phase. Under the canonical schedule with measured processing at $\alpha_S = 1$, no final-transport-outcome discordance is observed in the evaluated BCH and LDPC(512,256) TC-only or TM-only conditions, although some paired trials follow different canonical event paths. At some predefined TM-only processing-scale points, bidirectional final-outcome discordance is observed between success and P2-stage failure. The agreement state also varies with the evaluated finite contact time, and a schedule-phase case exhibits a larger transport-completion timing difference aligned with the TM slot structure. Under the evaluated conditions, the effect of omitting processing time therefore cannot be characterized by success or failure alone and must be interpreted together with event-path, timing, contact-time, and schedule conditions.

Keywords: Post-quantum cryptography · Over-the-air rekeying · CCSDS · Space Data Link Security · COP-1 · Processing time

1 Introduction

Long-lived space systems must update cryptographic keys during operation, and the transition to post-quantum cryptography (PQC) also affects this key-management process. The Space Data Link Security (SDLS) protocol of the Consultative Committee for Space Data Systems (CCSDS) provides a standardized framework for protecting space data links, while the SDLS Extended Procedures

support operational procedures including key management [5,2]. Hülising et al. proposed a PQC-based key-update mechanism for SDLS that sequentially exchanges three messages, P1, P2, and P3 [7]. When this exchange is transported over CCSDS Telecommand (TC) and Telemetry (TM) links, its behavior depends not only on cryptographic message sizes but also on the bidirectional transport and feedback procedures that carry those messages.

Transport-level evaluations may omit endpoint processing time and model each logical processing stage as completing instantaneously. However, TC Sequence-Controlled (Type-A) service uses Communications Operation Procedure-1 (COP-1) to control ordered delivery, retransmission, and recovery, while spacecraft-side reception state is returned to the ground through a Communications Link Control Word (CLCW) carried over the TM link [3,1,4]. Changing when a processing stage completes can therefore shift the relative timing among transmission, reception, feedback, timer, and recovery events. As a result, a zero-processing approximation can change the protocol execution path even when the final transport outcome remains unchanged.

This study evaluates that approximation by comparing a zero-processing baseline, T_0 , with a processing-aware model, T_1 , under identical CCSDS TC/TM and COP-1 semantics. The processing-aware model uses stage-level processing times measured from a study-specific Triple-KEM reference implementation, and predefined processing-scale factors are additionally applied to the spacecraft-side stages. The comparison separates three observation levels: final transport outcome, a canonical event path representing transport and control actions, and continuous timing. Finite contact time, direction-specific loss, and schedule phase are also evaluated so that agreement between T_0 and T_1 is interpreted together with the operating conditions under which it is observed.

This study makes three methodological contributions. First, it formulates processing-time omission as a paired model comparison in which transport, loss, COP-1, CLCW, and protocol semantics are held fixed and only the modeled processing time is changed. Second, it separates agreement in final transport outcome, canonical event path, and continuous timing so that execution differences hidden by outcome-only comparison can be observed explicitly. Third, it evaluates these agreement levels over predefined processing-scale, direction-specific loss, finite-contact, and schedule-phase conditions using the measured stage-level processing profile. This design is intended to identify the evaluated conditions under which the zero-processing approximation is observationally consistent with the processing-aware model, without interpreting the evaluation grid as a universal operational threshold.

The evaluation shows that agreement at the three observation levels does not always coincide. Under the canonical schedule with measured processing at $\alpha_S = 1$, no final-transport-outcome discordance is observed in the evaluated BCH and LDPC(512,256) TC-only or TM-only conditions, although some paired trials have different canonical event paths. At some predefined TM-only processing-scale points, final-outcome discordance is also observed, whereas selected higher-scale TC-only points contain many event-path disagreements but no final-outcome

discordance. The observed agreement state also varies with finite contact time and schedule phase. The effect of omitting processing time should therefore be described using final transport outcome, canonical event path, continuous timing, contact time, and schedule conditions rather than by a single success/failure result alone.

2 Background and Related Work

2.1 SDLS and Post-Quantum Key Update

The CCSDS Space Data Link Security (SDLS) protocol specifies security services for space data links by introducing security-related fields and procedures into data-link frames [5]. The SDLS Extended Procedures (SDLS-EP) provide additional procedures for SDLS operation and management, including key-management functions exchanged through protocol data units [2]. Together, these mechanisms provide a standardized basis for managing cryptographic keys over operational space links.

Hülsing et al. proposed a PQC-based SDLS key-update mechanism that sequentially exchanges three messages denoted P1, P2, and P3 [7]. In the final published version, the TK(Kyber512+X25519) configuration used as the transport input in this study has P1, P2, and P3 sizes of 1664, 1632, and 16 octets, respectively. The mechanism defines the cryptographic key-update exchange, but transporting those messages over a specific CCSDS TC/TM configuration additionally requires choices for encapsulation, fragmentation, channel coding, reliability control, and return-link feedback.

Accordingly, this study does not modify the cryptographic key-update protocol itself. Instead, it uses the published P1–P2–P3 exchange as the transport input, instantiates it over the CCSDS transport structure defined in Section 3, and analyzes the effect of endpoint processing time.

2.2 CCSDS TC/TM and COP-1

CCSDS Telecommand (TC) and Telemetry (TM) protocols provide the data-link structures used for command uplink and telemetry return traffic [3,4]. In the transport configuration of this study, P1 and P3 are carried in the TC direction, whereas P2 is returned in the TM direction. Because the two directions use different delivery and control structures, the P1–P2–P3 exchange cannot be represented as a single symmetric point-to-point transfer.

For TC Sequence-Controlled (Type-A) service, Communications Operation Procedure-1 (COP-1) controls ordered delivery, retransmission, and recovery through the interaction of the Frame Operation Procedure-1 (FOP-1) and the Frame Acceptance and Reporting Mechanism-1 (FARM-1) [1]. The spacecraft-side reception state is reported through a Communications Link Control Word (CLCW) carried on the TM return link. A TC recovery action can therefore depend not only on the uplink transmission history but also on the timing and delivery of return-link TM frames that carry CLCW information.

This temporal coupling is important when processing time is introduced into the model. Processing time can shift the readiness of subsequent P2 or P3 actions relative to TM slots, CLCW feedback, and COP-1 timer events without changing the underlying protocol rules. For this reason, this study distinguishes final transport outcome, canonical event path, and continuous timing when comparing the zero-processing and processing-aware models.

2.3 Related Work

Prior work has examined PQC in satellite communication from several perspectives. Mutlugun et al. evaluated a quantum-resistant IKEv2 design for satellite networking environments, providing an example of how PQC-based key-establishment mechanisms interact with satellite-network performance constraints [9]. Ghosh and Nath analyzed lattice-based cryptographic techniques for secure satellite communication using the performance characteristics of PQC primitives [6]. These studies address the broader interaction between PQC and satellite communication, but they do not model the P1–P2–P3 SDLS key-update exchange over the COP-1/CLCW-coupled CCSDS TC/TM structure considered here.

From an operational key-management perspective, Kim and Jang classified satellite OTAR technologies and summarized directions for subsequent research, emphasizing that satellite rekeying should be considered together with characteristics of the communication and operational environment [8]. Their work provides broader context for OTAR as an operational security function, whereas the present study focuses on a specific modeling question: how omitting processing time affects agreement between two otherwise identical CCSDS TC/TM execution models for PQC-based SDLS key-update transport.

The existing literature thus provides the cryptographic key-update mechanism, satellite-oriented PQC performance studies, and the broader OTAR context. Rather than proposing a new key-update protocol or PQC primitive, this study holds the key-update traffic and CCSDS transport semantics fixed and evaluates how well a zero-processing approximation preserves the final transport outcome, canonical event path, and continuous timing. The research question is therefore not whether PQC-based OTAR is universally feasible, but whether the zero-processing abstraction preserves the observations required by the analysis under the evaluated operating conditions.

3 System Model and Processing-Aware Implementation

3.1 Transport Model and Assumptions

This study models the sequential P1–P2–P3 exchange of the PQC-based key update described in Section 2.1 as a single end-to-end trial. The reference transport configuration uses TK(Kyber512+X25519) reported by Hülsing et al., with input sizes of 1664, 1632, and 16 B for P1, P2, and P3, respectively [7]. These

values are transport inputs for this study and do not imply that the same cryptographic implementation is used for the processing-time measurements described in Section 4.1. All times are expressed relative to $t = 0$, defined as the instant at which the communication window opens. For a finite contact time T_{contact} , contact close is defined at $t = T_{\text{contact}}$.

P1 and P3 are mapped to the TC uplink from the ground to the spacecraft, whereas P2 is mapped to the TM downlink from the spacecraft to the ground. P1 and P3 use TC Sequence-Controlled (Type-A) service, in which COP-1, consisting of FOP-1 and FARM-1, controls ordered frame delivery, retransmission, and recovery [3,1]. The state of spacecraft-side FARM-1 is returned to ground-side FOP-1 through the CLCW, which in this study is carried in the Operational Control Field (OCF) of a TM Transfer Frame [4,1]. Consequently, delivery of P1 and P3 involves not only forward TC transmission but also CLCW delivery over the TM return link and the resulting COP-1 state transitions and recovery actions.

P2 is carried in TM Transfer Frames and, in the transport configuration of this study, is delivered to the higher-level ground-side key-update procedure only after the complete P2 message has been reconstructed. When P2 data are not ready at a TM transmission opportunity, an idle frame may be used to preserve the continuous return-link slot structure and CLCW delivery. TM frames carrying P2 and idle TM frames are subject to the same return-link loss model, so loss of a TM frame prevents delivery of both its data and any CLCW carried in its OCF.

The TM return link is modeled as a continuous sequence of fixed-duration slots determined by the selected transport configuration. CLCW state is sampled according to a fixed nominal reporting period, and each sample is assigned to the first eligible TM slot whose start time is equal to or later than the sample time. Under the evaluated transport parameters, distinct CLCW samples do not share the same carrier slot. If the exchange reaches its natural terminal state such that the carrier slot corresponding to the final CLCW sample is never generated, that uncarried sample is excluded from the canonical event path. This slot and CLCW schedule is part of the transport instantiation defined in this study and does not imply that CCSDS specifies a single fixed CLCW schedule for a particular PQC key update.

Channel errors are abstracted as independent Bernoulli losses of complete post-decoding transmission units rather than as physical-layer bit errors. In the TC direction, one dispatched CLTU is the loss unit, whereas in the TM direction one TM Transfer Frame is the loss unit. The loss probabilities in the two directions are denoted by p_{TC} and p_{TM} , and the loss directions and probabilities used in the evaluation are defined in Section 4.3.

A finite contact time T_{contact} limits the interval during which a new RF link action may be initiated in a key-update trial. A new RF transmission may start only when its start time is strictly less than T_{contact} , while an already-started transmission whose reception occurs exactly at T_{contact} is considered delivered within the contact. Conversely, if a transmission starts before contact close

Table 1. Logical Processing Stages with Modeled Processing Time

No.	Stage	Location	Role
1	G_SEND1	Ground	P1 generation and pre-transmission processing
2	S_RECV1	Spacecraft	Security and protocol processing after P1 reception
3	S_SEND2	Spacecraft	P2 generation and pre-transmission processing
4	G_RECV2	Ground	Security and protocol processing after P2 reception
5	G_SEND3_PRE	Ground	Processing before P3 submission
6	G_SEND3_POST	Ground	Ground-side local processing after P3 submission
7	S_RECV3	Spacecraft	Final spacecraft-side processing after P3 reception

but its reception time after serialization and propagation exceeds T_{contact} , the corresponding transmission unit is not delivered to the receiving procedure. In addition to the finite-contact evaluation points, this study uses an **UNBOUNDED** reference for which no contact-time limit is applied.

Contact close terminates RF-link availability but does not forcibly terminate local processing. Ground- or spacecraft-side local processing that began before contact close, or that follows causally from an earlier event, may therefore continue after contact close. However, no new RF transmission, reception, or COP-1 recovery is initiated after contact close, and local processing completed after contact close does not retroactively change the transport outcome already determined for the finite contact.

3.2 Zero- and Processing-Aware Models

To isolate the effect of omitting processing time, we define a zero-processing baseline T_0 and a processing-aware model T_1 over identical transport and event semantics. Both models use the same protocol state transitions, data dependencies, TC/TM transmission rules, COP-1/CLCW processing rules, and contact-time rules. The directly varied inputs are the service times of the logical processing stages under comparison; those differences can in turn change not only subsequent event times but also event ordering and occurrence, retransmission and recovery behavior, and the final transport outcome.

The processing-time injection points are the seven logical stages listed in Table 1. The prefixes G and S identify the ground-side and spacecraft-side protocol locations of the stages; they do not imply that the measurements were taken on two physically distinct systems.

Let $\hat{d}_i$ denote the median processing time measured for logical stage i in the reference implementation described in Section 4.1. The main evaluation uses this fixed median for each stage, while per-trial resampling of processing time is treated separately as a sensitivity analysis.

In the zero-processing baseline T_0 , the seven processing events and their data dependencies are retained, but the service time of every compared stage is set to zero.

$$d_i^{(T_0)} = 0 \quad \forall i. \quad (1)$$

In the processing-aware model T_1 , the measured median is used directly for ground-side stages and is multiplied by the spacecraft processing-scale factor α_S for spacecraft-side stages.

$$d_i^{(T_1)} = \begin{cases} \hat{d}_i, & i \in \mathcal{G}, \\ \alpha_S \hat{d}_i, & i \in \mathcal{S}, \end{cases} \quad (2)$$

where

$$\mathcal{G} = \{\text{G_SEND1}, \text{G_RCV2}, \text{G_SEND3_PRE}, \text{G_SEND3_POST}\}$$

and

$$\mathcal{S} = \{\text{S_RCV1}, \text{S_SEND2}, \text{S_RCV3}\}.$$

The condition $\alpha_S = 1$ applies the measured spacecraft-side processing times without additional scaling. Because ground-side processing remains present when $\alpha_S = 0$, $T_1(\alpha_S = 0)$ is not identical to T_0 . Moreover, α_S is an analytical parameter that scales the measured spacecraft-side processing vector; it is not a measured performance ratio or CPU-speed ratio for a particular spacecraft processor.

Even in T_0 , where service times are zero, the logical boundaries corresponding to enqueue, completion, and the resulting protocol effects of each processing stage are retained. Thus, T_0 is not a separate protocol that removes processing stages; it uses the same event-processing rules and data dependencies as T_1 while assigning zero service time to the selected stages. A zero-service-time stage still completes and triggers its subsequent effects at the same logical time, with ordering relative to other same-time events determined by the fixed event-priority rules of the model.

3.3 Agreement Metrics and Comparison Rules

Differences between T_0 and T_1 are separated into three observation layers: final transport outcome, canonical event path, and continuous timing. These layers capture different levels of execution information, so equality of the final transport outcome is not taken to imply equality of the event path or the major completion times.

Transport Outcome For a finite contact time T , a stage is considered completed within the contact only if its completion event exists and its completion time is no later than T . For P1 and P2, $t_{\text{P1,secure}}$ and $t_{\text{P2,secure}}$ denote the times at which receiver-side security processing and delivery to the higher-level key-update procedure are completed. The P3 transport-completion time $t_{\text{transport_complete}}$

is the time at which the last required P3 TC frame is successfully accepted by spacecraft-side FARM-1 and the P3 transport-completion condition becomes satisfied.

The stage-completion conditions for finite contact time T are defined as follows.

$$\begin{aligned} C_1(T) &\equiv [t_{P1,secure} \text{ defined}] \wedge t_{P1,secure} \leq T, \\ C_2(T) &\equiv C_1(T) \wedge [t_{P2,secure} \text{ defined}] \wedge t_{P2,secure} \leq T, \\ C_3(T) &\equiv C_1(T) \wedge C_2(T) \wedge [t_{transport_complete} \text{ defined}] \wedge t_{transport_complete} \leq T. \end{aligned} \quad (3)$$

At the contact-unbounded reference point **UNBOUNDED**, no contact-time upper bound is applied, and completion is determined only by whether the corresponding completion event occurs in the actual execution.

$$\begin{aligned} C_1^{UB} &\equiv [t_{P1,secure} \text{ defined}], \\ C_2^{UB} &\equiv C_1^{UB} \wedge [t_{P2,secure} \text{ defined}], \\ C_3^{UB} &\equiv C_1^{UB} \wedge C_2^{UB} \wedge [t_{transport_complete} \text{ defined}]. \end{aligned} \quad (4)$$

Even if local processing completes after contact close and additional processing timestamps are recorded, the finite-contact completion classification in Eq. (3) is not changed retroactively. The final transport outcome under a given contact time and local-processing timestamps that may continue to be recorded after contact close are therefore treated as distinct observables.

The final transport outcome Q of each trial belongs to one of four categories.

$$Q \in \{\text{SUCCESS}, \text{FP1}, \text{FP2}, \text{FP3}\}. \quad (5)$$

FP1 denotes a trial in which C_1 is not satisfied, **FP2** a trial in which C_1 is satisfied but C_2 is not, and **FP3** a trial in which C_1 and C_2 are satisfied but C_3 is not. **SUCCESS** denotes a trial in which C_3 holds. These four categories are mutually exclusive and collectively exhaustive and are distinguished from individual diagnostic events such as frame loss, COP-1 T1 timer expiration, retransmission, or Alert. Because a trial may still satisfy the completion conditions through subsequent recovery after such a diagnostic event occurs, an individual diagnostic event is not itself interpreted as a stage failure.

Canonical Event Path To compare protocol execution in addition to the final transport outcome, we define the canonical event path Π . The canonical event path is a token sequence constructed by ordering the event records generated by the execution or finite-contact projection in time, removing the numerical timestamps themselves, and retaining the semantic event types, identifiers, and outcomes. Numerical timestamps are therefore excluded from equality comparison, while absolute TM slot indices and the CLCW sample and carrier-slot identifiers that distinguish schedule position are retained. Two canonical event paths are

Table 2. Main Information Retained in the Canonical Event Path

Category	Information used for equality comparison
TC	P1/P3 identity, logical frame, transmission attempt, transmission and reception outcomes, and FARM-1 acceptance or discard with discard classification
TM	Absolute TM slot index, P2 or idle-frame identity, P2 logical frame, slot outcome, and whether OCF and CLCW are carried
CLCW	Sample index, reported COP-1 state, carrier-slot index, and carrier-frame outcome
COP-1	FOP-1/FARM-1 state transitions, COP-1 T1 timer expiration, recovery decisions, and Alerts
Protocol	P1/P2 secure delivery, P3 submission, P3 transport completion, and final transport outcome

considered equal only when the token sequences have the same length and every field used for equality comparison matches at each position.

Table 2 summarizes the main information retained in the canonical event path.

Conversely, CPU-stage enqueue, start, and completion events, processing service times, queueing delays, numerical timestamps, and raw random values are excluded from canonical-path equality comparison. Processing-time observation points such as key-ready completion are also excluded from the canonical event path itself and are instead treated as continuous timing information.

The canonical event path terminates with a token representing the final transport outcome Q of the trial. Canonical-path equality therefore implies equality of the final transport outcome.

$$\Pi_{T_0} = \Pi_{T_1} \implies Q_{T_0} = Q_{T_1}. \quad (6)$$

The converse does not hold in general. Two trials may reach the same SUCCESS, FP1, FP2, or FP3 outcome while following different TC/TM transmissions, CLCW feedback, or COP-1 recovery events and therefore different canonical event paths.

Paired Agreement Regimes Paired T_0 and T_1 executions under the same experimental condition are classified as $R0$, $R1$, or $R2$ according to the relationship between their final transport outcomes and canonical event paths.

$$\begin{aligned} R0 : & \quad Q_{T_0} = Q_{T_1} \wedge \Pi_{T_0} = \Pi_{T_1}, \\ R1 : & \quad Q_{T_0} = Q_{T_1} \wedge \Pi_{T_0} \neq \Pi_{T_1}, \\ R2 : & \quad Q_{T_0} \neq Q_{T_1}. \end{aligned} \quad (7)$$

The three states are mutually exclusive and collectively exhaustive over paired trials. By Eq. (6), an $R2$ pair cannot have equal canonical event paths, so $R2$ is not subdivided further by path equality. An $R0$ pair, however, only means that the final transport outcome and canonical event path are equal; it does not

Table 3. Main Timing Observation Points for Continuous Timing Comparison

Observation point	Definition
$t_{\text{transport_complete}}$	Time at which all required P3 TC frames have been accepted by FARM-1 and the transport-completion condition is satisfied
$t_{\text{ground_key_ready}}$	Completion time of G_SEND3_POST processing
$t_{\text{spacecraft_key_ready}}$	Completion time of S_RECV3 processing
$t_{\text{protocol_complete}}$	Maximum of the ground and spacecraft key-ready times when both exist

require every event to occur at the same absolute time or every timing observation point to be equal.

The labels $R0$, $R1$, and $R2$ classify individual paired trials. For stochastic experimental conditions containing multiple trials, an experimental point is not forced into a single regime; instead, we report the composition of $R0$, $R1$, and $R2$ trials observed at that point. Trial counts and aggregation procedures are described in Section 4.4.

Continuous Timing Because canonical-path equality excludes numerical timestamps, residual timing differences between the two models are evaluated separately as continuous quantities. The main timing observation points used in this study are listed in Table 3.

The protocol-completion time is defined only when both key-ready times exist.

$$t_{\text{protocol_complete}} = \max(t_{\text{ground_key_ready}}, t_{\text{spacecraft_key_ready}}). \quad (8)$$

Ground- or spacecraft-side key-ready times may be recorded later than T_{contact} as a result of local processing that is allowed to continue after contact close. Timing observation points in finite-contact projections are therefore not forcibly truncated at the contact time. In contrast, $t_{\text{transport_complete}}$ requires actual link delivery and FARM-1 acceptance, so when it is defined under finite contact it necessarily lies within the contact interval.

For timing observation point x , the signed timing difference is defined as the processing-aware value minus the zero-processing-baseline value.

$$\Delta t_x = t_{x,T_1} - t_{x,T_0}. \quad (9)$$

Δt_x is defined only for paired trials in which the corresponding observation point exists in both T_0 and T_1 . If the timestamp is absent from either execution, it remains undefined rather than being replaced by zero, infinity, the contact time, or another substitute value. In the analysis, all paired trials for which Δt_x is defined are aggregated separately for each observation point and experimental condition; success, $R0$, and event-path equality are not imposed as additional

eligibility conditions. The support and summary statistics for each aggregation are described in Section 4.4.

The schedule-phase sensitivity additionally uses a binary success/failure outcome separately from the four-class final transport outcome. This binary comparison is an auxiliary measure for observing whether success status differs under a particular phase condition and is not interpreted as the same metric as $R2$, which is defined from the four-class final transport outcome.

3.4 Implementation Overview

The implementation processes the processing stages, TC/TM transmissions, COP-1/CLCW operations, and protocol-state transitions defined above within a single deterministic discrete-event structure. T_0 and T_1 are not separate protocol implementations; they share the same event-handling rules, and the directly varied inputs are the processing service times defined in Section 3.2.

Because multiple events may occur at the same logical time, the implementation uses a fixed event-priority order. Reception and the resulting FARM-1/FOP-1 state updates are applied before subsequent recovery decisions at the same time. A CLCW arrival is processed before a simultaneous COP-1 T1 timer expiration, while contact close exactly at T_{contact} takes precedence over a timer expiration and the initiation of a new RF transmission at the same time. The canonical event path is likewise constructed by ordering events according to this fixed order and then removing numerical timestamps.

For a finite contact time, the outcome assigned to a transmission unit is determined by evaluating the contact boundary before channel loss. Let t_s be the scheduled start time of a transmission according to the link schedule, t_a the scheduled reception time if the transmission starts, and T the contact time. An RF transmission actually starts only when $t_s < T$, and the outcome is classified in the following priority order.

$$O(t_s, t_a, \ell, T) = \begin{cases} \text{NOT_TRANSMITTED_CONTACT_CLOSE}, & t_s \geq T, \\ \text{NOT_DELIVERED_CONTACT_CLOSE}, & t_s < T \wedge t_a > T, \\ \text{LOST}, & t_s < T \wedge t_a \leq T \wedge \ell = 1, \\ \text{DELIVERED}, & t_s < T \wedge t_a \leq T \wedge \ell = 0. \end{cases} \quad (10)$$

Here, ℓ indicates whether the corresponding transmission unit is lost by the channel. A transmission that starts before contact close but would be received after contact close is therefore classified as `NOT_DELIVERED_CONTACT_CLOSE` regardless of the loss variate, while a reception exactly at T is delivered if no channel loss occurs. A transmission whose scheduled start is at or after the contact time is never initiated as an RF transmission and is recorded as `NOT_TRANSMITTED_CONTACT_CLOSE`.

The evaluation first generates a contact-unbounded execution until its natural terminal state and then constructs each finite-contact result by projecting the contact boundary onto that execution. The projection is not a simple truncation of

the event record at time T ; it re-evaluates transmission-unit delivery at the contact boundary and propagates the resulting event dependencies and permitted post-close local-processing chains to construct the final transport outcome, canonical event path, and timing observations. A natural terminal state is reached when no further local processing or recoverable TC action remains and the P1–P2–P3 exchange has reached either success or a defined failure state. If natural termination occurs before a particular finite contact time, no additional periodic idle TM frames or CLCW events are generated after that point.

When finite contact is projected, no new RF transmission, reception, recovery driven by CLCW return, or timer action is processed after contact close. However, local processing that follows causally from the protocol state before contact close continues, and a completed local stage may initiate another local-processing stage. For example, this chain may include `G_SEND3_PRE`, P3 submission, and `G_SEND3_POST` after `G_RECV2`. Thus, if post-close local processing emits a P3-submission event, the corresponding `P3_SUBMIT` protocol token may appear in the canonical event path, but no new P3 RF transmission can begin afterward. CPU enqueue, start, and completion events and key-ready completion itself are recorded as timing information rather than as canonical-path tokens.

Each CLCW sample is associated with the first eligible TM slot whose start time is equal to or later than the sample time. Under the evaluated transport parameters, TM slots are sufficiently frequent relative to the CLCW reporting period that distinct samples do not contend for the same carrier slot. If the carrier slot for the final CLCW sample generated shortly before natural termination lies after the natural terminal state and the carrier frame is therefore never generated, that final uncarried sample is excluded from the canonical event path. By contrast, if contact close occurs before natural termination and prevents the associated CLCW carrier slot from actually starting, the blocked carrier is represented as `NOT_TRANSMITTED_CONTACT_CLOSE`. The final CLCW carrier suppressed by natural termination and a CLCW carrier blocked by contact close are therefore handled by different rules. For other normally generated CLCW tokens, the sample index, reported state, carrier slot, and carrier-frame outcome are retained.

Probabilistic paired comparisons use common random numbers (CRN) to couple T_0 and T_1 through deterministic random-variable identifiers. TC-loss variates are identified by trial number, protocol stage, logical frame number, and transmission-attempt number, whereas TM-loss variates are identified by trial number and absolute TM slot index. The same random value is used for the same identifier, rather than consuming random values sequentially according to transmission order. Because processing-time differences can change the sets and times of transmissions and slots actually referenced by the two models, CRN coupling is not interpreted as giving the models identical physical transmission histories or identical loss traces. The experimental-point construction, trial counts, processing scales, contact-time evaluation points, and analysis aggregation procedures are described in Section 4.

4 Experimental Methodology

4.1 Processing Measurement

To construct the processing-time inputs for the processing-aware model T_1 , we measured the execution times of the seven logical processing stages defined in Section 3.2. The transport model uses the P1, P2, and P3 message sizes reported by Hülsing et al. for TK(Kyber512+X25519) as transport inputs [7], whereas processing time is measured using a separate Triple-KEM reference implementation built for this study. The reference implementation combines ML-KEM-512 and DHKEM(X25519) using OpenSSL 3.5.8 and is not the original Kyber512 implementation used by Hülsing et al. The measured values therefore constitute a study-specific processing profile used together with the selected transport inputs; they do not represent the performance of the original authors' implementation or of an operational spacecraft processor.

The measurement target is not an isolated cryptographic primitive but the seven study-defined processing stages executed within a normal Triple-KEM handshake. Within each timed session, the processing times of G_SEND1, S_RECV1, S_SEND2, G_RECV2, G_SEND3_PRE, G_SEND3_POST, and S_RECV3 are measured in one complete handshake. The 5,000 samples are therefore 5,000 consecutive complete-handshake samples from one measurement run rather than 5,000 independent benchmark campaigns.

The stage boundaries include the hybrid key generation, encapsulation or decapsulation, authenticated encryption or decryption, transcript hashing, confirmation processing, and final key derivation performed within the corresponding stage. They also include parsing, serialization, copying, state updates, per-operation EVP/provider setup, and use of the initialized random generator performed inside that stage. The ephemeral public-key import performed in S_RECV1 is likewise included in that stage. Process startup, provider loading, long-term key generation and provisioning, session creation and destruction, file I/O, SDLS frame protection, CCSDS framing, COP-1 processing, RF-link behavior, and P3 link-layer submission are excluded from the measurement boundary. The resulting measurements are therefore integrated stage-level service times injected into T_1 , not isolated benchmarks of ML-KEM-512, X25519, or other individual cryptographic primitives.

Table 4 summarizes the measurement environment and execution configuration.

The measurement run was pinned to one logical CPU, but CPU frequency was not fixed; the governor and turbo state were recorded as part of the measurement environment. The reported 1 ns value is the operating system's clock resolution and is not interpreted as a guarantee of 1 ns timing accuracy for individual processing measurements. The resulting measurements therefore describe a processing profile conditional on the reference environment in Table 4.

After 256 warm-up handshakes, 5,000 complete handshakes were collected as timed samples. For each stage i , the predefined analysis procedure computes the sample count, minimum, 5th, 25th, 50th, 75th, 95th, and 99th percentiles,

Table 4. Reference Configuration for Processing-Time Measurement

Item	Configuration
Platform	Intel Core i7-7740X, Ubuntu 24.04.4 LTS, x86-64
Cryptographic implementation	OpenSSL 3.5.8 default provider, ML-KEM-512 + DHKEM(X25519), AES-256-GCM, SHA3-256, HMAC-SHA256
Measurement-harness build	GCC 13.3.0, <code>-std=c11 -O3 -DNDEBUG -D_POSIX_C_SOURCE=200809L</code>
OpenSSL build	<code>linux-x86_64 -release no-shared</code> , static <code>libcrypto.a</code> , upstream release optimization
Additional optimization	No <code>-march=native</code> , LTO, PGO, sanitizer, or profiler
CPU execution condition	Pinned to one logical CPU (CPU4); <code>powersave</code> governor, turbo enabled, no fixed frequency
Warm-up / measurement	256 / 5,000 complete handshakes
Timer	<code>CLOCK_MONOTONIC_RAW</code> ; reported clock resolution 1 ns
Post-processing	No outlier removal; no timer-overhead subtraction

maximum, mean, sample standard deviation, and median absolute deviation. Quantiles are computed by sorting the observations and linearly interpolating between adjacent values at position $(n - 1)q$. The main T_1 evaluation uses the median $\hat{d}_i$ for each stage. When the processing-time distribution is visualized in Section 5.1, the stored 5th and 95th percentiles from the same summary are shown together with the median. The p05–p95 interval summarizes the empirical stage-latency distribution and is not a confidence interval for the median or a statistical-significance interval.

4.2 Component-Level Reconstruction Cross-Check

To support interpretation of the integrated stage measurements A' , stage-level processing times are reconstructed from medians of component operations measured separately using the same cryptographic library and implementation family. This reconstruction, denoted B , is not an independent implementation validation of A' ; it is a supporting cross-check of how each integrated stage can be described as a combination of separately measurable component operations.

The component benchmark comprises 18 operations, including hybrid key generation, encapsulation and decapsulation, AEAD processing for different input lengths, transcript hashing, protocol-state processing, and final key derivation. Each measurement cycle executes each of the 18 operations once, and a cyclic rotation of the starting position is applied across cycles so that the same operation does not always occupy the same order position. The reconstruction uses component medians obtained from 5,000 measurement cycles following 256 warm-up cycles per component.

The component-based reconstruction for stage i is expressed as

$$B_i = \sum_{j=1}^M m_{ij} \tilde{d}_j, \quad M = 18, \quad (11)$$

where $\tilde{d}_j$ is the median of separately measured component j , and m_{ij} is the number of occurrences of that operation in stage i under the predefined reconstruction rule.

The difference between the integrated measurement and component reconstruction is defined as the signed reconstruction residual

$$r_i = \hat{d}_i^{A'} - B_i. \quad (12)$$

Although A' and B use the same software family, their measurement boundaries and execution contexts are not identical. Because A' measures an entire stage within a normal complete handshake, it naturally includes parsing, serialization, copying, state checks and updates, inter-operation glue processing, and other work inside the stage boundary. By contrast, B combines separately measured component medians according to Eq. (11) and does not force all such integrated-stage work to be included. Moreover, the sum of component medians and the median of an integrated stage are not the same statistical object, and their execution contexts and locality can differ.

Accordingly, r_i is not interpreted as an estimate of a specific integration overhead. The sign or magnitude of the residual is not used as causal evidence for a particular implementation element, optimization effect, or measurement error. Likewise, whether B_i lies inside the p05–p95 interval of A' is not used as a statistical validation or significance test of the reconstruction. The study uses A' as the reference processing profile applied to T_1 , while B and r_i are used only as supporting descriptions of agreement and discrepancy in stage composition.

4.3 Experimental Conditions and Scenario Design

Reference Transport Configuration The evaluation performs paired comparisons of T_0 and T_1 over the same CCSDS TC/TM and COP-1 semantics defined in Section 3. The reference key-update traffic is TK(Kyber512+X25519), and BCH and LDPC(512,256) are evaluated separately as TC channel codes. The reference TM return link uses the Turbo rate-1/3 profile with $L_{\text{TM}} = 446$ octets, giving a TM slot serialization time of $T_{\text{slot}} = 5.406$ ms. The TC and TM coded-symbol rates are 10 ksymbol/s and 2 Msymbol/s, respectively; the one-way propagation delay τ_{prop} is 6.44 ms, and the CLCW reporting period T_{CLCW} is 1 s.

COP-1 uses FOP sliding-window width $K = 6$, FARM window $W = 12$, $PW = NW = 6$, Transmission Limit 3, and Timeout Type 0. No additional margin is added to Timer Initial; instead, the following lower bound is used.

$$T_{1,\min} = t_{\text{proc,other}} + t_b + 2\tau_{\text{prop}} + g_{\max} + T_{\text{slot}}, \quad (13)$$

where $t_{\text{proc,other}} = 0$ and $T_{1,\text{Initial}} = T_{1,\min}$. Here, t_b is the serialization time of the CLTU generated from a maximum-length TC Transfer Frame under the

selected TC channel code, excluding the PLOP-2 Idle Sequence. The term $g_{\max}$ is the maximum interval between the start times of consecutive realized CLCW-carrier slots in the corresponding configuration. When schedule phase or transport configuration changes, $g_{\max}$ and $T_{1,\min}$ are recomputed for that configuration. Consequently, the phase sensitivities described below are not interpreted as pure phase shifts under a fixed timer value. For a fixed transport configuration, changing the finite contact time T_{contact} does not itself change $T_{1,\min}$.

Loss-Level Selection The stochastic loss probability p^* is not selected as an estimate of a representative satellite-channel loss probability. Instead, a predefined pilot-selection rule is applied to a candidate loss grid to choose one common loss level for the subsequent stochastic evaluation. The pilot-selection rule is fixed before the pilot results are inspected.

The candidate set is

$$\mathcal{P}_{\text{pilot}} = \left\{ \frac{1}{10}, \frac{1}{5}, \frac{2}{5} \right\}. \quad (14)$$

For each candidate, BCH and LDPC(512,256) are evaluated under both TC-only and TM-only loss with the canonical schedule. For code c , loss direction d , and candidate p , let $q_{c,d}(p)$ be the UNBOUNDED terminal success fraction of the pilot reference processing-aware condition $T_{1,\text{ref}}$. Each pilot cell contains 512 paired trials.

The selection score is defined as

$$V_{c,d}(p) = q_{c,d}(p) (1 - q_{c,d}(p)), \quad S(p) = \min_{c,d} V_{c,d}(p). \quad (15)$$

The rule selects the candidate with the largest $S(p)$, with an exact tie broken in favor of the smaller p . Applying this predefined rule to the pilot results selects $p^* = 2/5$. The predefined saturation fallback is not invoked in this pilot. Thus, $p^* = 2/5$ is a study-specific stochastic stress condition selected from pilot results by a predefined rule; it is not a representative physical channel quality or a universal loss threshold for satellite links.

The TC-only condition in the final evaluation is defined as

$$(p_{\text{TC}}, p_{\text{TM}}) = (p^*, 0) \quad (16)$$

and the TM-only condition as

$$(p_{\text{TC}}, p_{\text{TM}}) = (0, p^*) \quad (17)$$

respectively. No stochastic condition with simultaneous TC and TM loss is included in the final scenario set.

Processing-Scale Grid The code-specific reference points for the processing-scale grid are obtained using a deterministic no-loss landmark-search procedure fixed in advance. The search procedure is therefore fixed before stochastic production evaluation, whereas the numerical landmark values themselves are results of the

Table 5. Scenario Families and Primary Analysis Scope

Family	No.	Trials N	Varied axis	Primary observables
CORE_TO_VS_T1_ALPHA	48	No-loss 1 / loss 4096	α_S	4-class outcome, $R0/R1/R2$, path, timing
SENS_TM_PHASE	48	No-loss 1 / loss 4096	TM phase	path, binary outcome, timing
SENS_CLCW_PHASE	48	No-loss 1 / loss 4096	CLCW phase	path, binary outcome, timing
SENS_PROCESSING_RESAMPLING	6	All 4096	processing sample	path, binary outcome, timing
SENS_P1_READY_AT_OPEN	6	No-loss 1 / loss 4096	contact-open readiness	path, binary outcome, timing
Total	156	434,226 paired trials	CORE 48 + SENS 108	—

deterministic search. Under the canonical schedule, the search uses $T_1(\alpha_S = 0)$ as the structural reference and identifies a code-specific reference point at which structural behavior changes as the spacecraft-side processing scale increases. This reference is not the zero-processing model T_0 , because ground-side processing remains present in $T_1(\alpha_S = 0)$.

The code-specific landmarks obtained from the search are

$$L_{\text{BCH}} = \frac{6427303555}{4294967296} \approx 1.496473, \quad (18)$$

$$L_{\text{LDPC}} = \frac{3965698833}{536870912} \approx 7.386690. \quad (19)$$

The final production processing-scale set for code c is constructed as

$$\mathcal{A}(c) = \text{sort unique} \left(\{0, 1\} \cup \left\{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1, \frac{3}{2}, 2\right\} L_c \right). \quad (20)$$

Duplicates are removed using exact rational representations, yielding eight α_S evaluation points for both BCH and LDPC(512,256). The multiplier rule and final-grid construction in Eq. (20) are fixed before the stochastic processing-scale results are inspected. No interpolation is performed between evaluation points, and neither a landmark nor its multiples are interpreted as continuous processing-time thresholds. The parameter α_S is an analytical scaling factor for the spacecraft-side processing vector defined in Section 3.2, not a measured performance ratio of an actual spacecraft CPU.

Scenario Families and Sensitivity Axes The final evaluation campaign comprises the 156 scenarios in Table 5, totaling 434,226 paired trials. The CORE family is the primary evaluation over α_S , while the remaining 108 scenarios are sensitivity families that vary one processing or schedule assumption at $\alpha_S = 1$. Each sensitivity axis is evaluated one factor at a time; sensitivity axes are not crossed with one another or factorially combined with the processing-scale axis.

For the TM-phase sensitivity,

$$\phi_{\text{TM},k} = \frac{kT_{\text{slot}}}{8}, \quad k = 0, \dots, 7 \quad (21)$$

varies the TM slot-grid phase while the CLCW phase is fixed at canonical $k = 0$. For the CLCW-phase sensitivity,

$$\phi_{\text{CLCW},k} = \frac{kT_{\text{CLCW}}}{8}, \quad k = 0, \dots, 7 \quad (22)$$

varies the CLCW reporting phase while the TM phase is fixed at canonical $k = 0$. Because $T_{\text{CLCW}} = 1$ s, the evaluated CLCW phases are 0, 0.125, 0.250, 0.375, 0.500, 0.625, 0.750, and 0.875 s. When the CLCW phase is zero, the first nominal sample occurs at T_{CLCW} ; for a positive phase, the first sample occurs at $\phi_{\text{CLCW},k}$ and subsequent samples repeat every T_{CLCW} . The two phase families are separate eight-point one-factor sensitivities; the 8×8 Cartesian product of TM phase and CLCW phase is not evaluated. For each phase configuration, g_{max} and $T_{1,\text{Initial}}$ are recomputed from the modified schedule geometry.

The processing-resampling sensitivity replaces the fixed-median processing profile in T_1 with endpoint-vector draws from the empirical samples in Section 4.1. On the ground side, the four values for `G_SEND1`, `G_RECV2`, `G_SEND3_PRE`, and `G_SEND3_POST` are drawn together from the same measurement row, while the three values for `S_RECV1`, `S_SEND2`, and `S_RECV3` are drawn together from the same spacecraft-side measurement row. The ground and spacecraft vectors use separate deterministic keyed draws. Thus, empirical joint structure is preserved within each endpoint vector, but the full seven-stage dependence between ground and spacecraft values from an original complete-handshake sample is not preserved. Sampling is with replacement, and the same trial ID and endpoint deterministically map to the same sample index. This family uses $\alpha_S = 1$, while T_0 retains the original zero-service model. Because the processing vector itself is stochastically resampled, all six processing-resampling scenarios use $N = 4096$, including the no-loss scenarios. This sensitivity evaluates variability in the measured empirical endpoint vectors and is not interpreted as a general robustness analysis for arbitrary hardware-processing variation or actual spacecraft CPUs.

The P1-ready-at-open sensitivity modifies the baseline `STRICT_ONLINE` assumption under which `G_SEND1` is enqueued after contact open. In this sensitivity, fresh single-use P1 bytes and the matching ground session and protocol state are assumed to be ready at contact open. The complete `G_SEND1` computation is therefore moved outside the contact-critical timeline rather than removed, and its processing cost remains identified as a pre-contact cost. After setting $t_{\text{P1,protocol_ready}} = 0$, the prepared P1 is submitted to the existing TC scheduler from $t = 0$; the subsequent RF transmission, COP-1, and loss semantics are unchanged. Both T_0 and T_1 use the same ready-at-open policy. This sensitivity therefore does not evaluate $\alpha_G = 0$, long-term precomputation, or reuse of P1 across sessions.

Finite-Contact Projection Finite-contact projection is applied only to the 48 CORE scenarios and not to the 108 sensitivity scenarios. Each CORE execution

is first generated to its contact-unbounded natural terminal state and is then projected to the following 16 predefined horizons according to the semantics in Section 3.4.

$$\mathcal{H} = \{1.0, 1.8, 2.4, 3.1, 4.1, 5.4, 7.2, 9.6, 12.7, 16.9, 29.7, 52.2, 91.8, 161.5, 284.2, 500.0\} \text{ s.} \quad (23)$$

UNBOUNDED is a separate reference that is not an element of Eq. (23) and is not treated as equivalent to 500 s. The 16 finite-contact results for a given scenario are projections of the same underlying paired execution rather than 16 independent Monte Carlo runs. Results at different horizons are therefore not treated as independent samples, no interpolation is performed between adjacent horizons, and no continuous contact-time threshold is inferred. Section 5.4 reports the BCH and LDPC(512,256) TM-only landmark cells in detail; this is a reporting selection from the existing CORE scenarios rather than an additional experimental point.

4.4 Paired Execution and Analysis Protocol

Paired Execution and Common Random Numbers For each scenario, T_0 and T_1 form a paired execution using the same transport parameters and common random numbers (CRN). A given random identifier maps to the same random value in both models, but processing-time differences can change the set of transmissions, TM slots, and random identifiers that are actually referenced. CRN therefore couples corresponding random variables but does not force the two models to follow identical physical transmission histories or identical loss traces.

The three observation layers defined in Section 3.3 are final transport outcome Q , canonical event path Π , and continuous timing Δt_x . $R0$, $R1$, and $R2$ are paired classifications constructed from Q and Π , not the three observation layers themselves. In particular, $R0$ does not mean that both models succeed; it means that they have the same four-class outcome and the same canonical event path. A pair in which both models have the same failure outcome and path can therefore also be classified as $R0$.

CORE Aggregation Each paired trial in the CORE family is classified as $R0$, $R1$, or $R2$ according to Section 3.3. For stochastic CORE cells, we compute the count and proportion of each class and use the stored pointwise Wilson 95% confidence interval for each proportion [10]. Each interval is pointwise for the corresponding scenario and class proportion and is not a simultaneous confidence region for the full α_S curve or parameter grid. An observed $R2$ count of zero in a stochastic cell therefore does not prove that the population discordance probability is exactly zero.

Deterministic no-loss CORE executions are treated as $N = 1$ controls. For these cells, the CORE composition retains the count and deterministic label, but the proportion and Wilson interval are not reported as probability estimates.

Deterministic results are not interpreted as stochastic probability estimates or sampling uncertainty.

For stochastic CORE cells containing $R2$ trials, directional macro transitions among SUCCESS, FP1, FP2, and FP3 are retained separately in the $T_0 \rightarrow T_1$ direction. Outcome discordance is therefore not reduced to a single total; the source and destination outcome classes are distinguished.

Continuous Timing Aggregation For each timing endpoint x , we use the signed difference defined in Eq. (9),

$$\Delta t_x = t_{x,T_1} - t_{x,T_0}$$

for the paired analysis. Δt_x is defined only when the endpoint exists in both T_0 and T_1 for the same paired trial. If the endpoint is absent from either model, the pair is excluded from that endpoint’s timing aggregation rather than assigning zero, the contact time, infinity, or another substitute value. Success status, $R0$, and canonical-path equality are not imposed as additional eligibility conditions for timing support.

For each scenario and endpoint, the number of eligible pairs is recorded as the support, and the median, 5th percentile, and 95th percentile of Δt_x are computed. If no eligible pair exists, support is zero and the median and quantiles remain undefined. The analysis output represents these undefined statistics as null; this rule does not imply that support-zero cases occur among the timing cells reported in this paper.

Timing quantiles are computed using a fixed rule that sorts exact-rational Δt_x values and linearly interpolates between adjacent values at position $(n - 1)q$. This is a separately defined timing-analysis contract rather than an automatic reuse of the processing-measurement summary rule from Section 4.1. For deterministic timing with $N = 1$, the median, q05, and q95 are the same single value, but this is not interpreted as an estimate of a stochastic distribution.

The timing summaries are stored in ns. The μs and ms values displayed in the paper are obtained by unit conversion only after timing aggregation is complete, and display rounding is likewise applied only after aggregation. Rounded display values are therefore not used for scenario joins, support determination, quantile computation, or any subsequent analysis.

Sensitivity Analysis The CORE-specific four-class $R0/R1/R2$ classification is not automatically assigned to SENS families. TM-phase, CLCW-phase, processing-resampling, and P1-ready-at-open sensitivities use canonical-event-path equality, paired binary success/failure outcomes, and continuous timing as their primary observables. For the binary outcome, the two discordant directions— T_0 success/ T_1 failure and T_0 failure/ T_1 success—are retained separately. This binary comparison is an auxiliary collapse of the four-class transport outcome into two categories and is not the same metric as CORE $R2$. Binary agreement in a sensitivity family therefore does not by itself establish four-class transport-outcome equivalence.

When stochastic sensitivity cells report the T_1 success probability, event-path equivalence, or either binary discordance proportion, the corresponding stored

Table 6. Validation Layers and Their Evidential Scope

Target	Property and scope checked	Result	Interpretation limit
Processing-protocol implementation	Normal P1–P2–P3 exchange, stage-level protocol state and outputs, and final-key agreement at both endpoints in the study-specific Triple-KEM	Passed the predefined correctness-validation procedure	Does not establish equivalence to the original authors’ implementation or provide a cryptographic security proof
Measurement harness	Complete handshake, final-key agreement, expected message structure, and zero timer calls in the untimed smoke branch of the actual measurement harness	Untimed functional smoke passed	Does not independently validate timing accuracy; timed runs place timers around the same stage calls
Final CORE replay	Comparison of 22 raw fields between final raw records and same-source replay for the 48 CORE scenarios under final SNAP-009	131,088/131,088 paired records exact match	Establishes same-source deterministic reproducibility, not equivalence to an independent implementation
Finite-contact projector	Comparison of 16-horizon projections with direct finite executions for selected CORE pairs under SNAP-006	16,896/16,896 equivalent	Applies to SNAP-006 and is not a full revalidation of the final SNAP-009 projector
Post-correction regression	Targeted direct-versus-projected comparison and related regression/preflight for boundary cases affected by uncarried CLCW and contact-close same-time ordering corrections	Targeted comparisons and regressions passed	Targeted evidence for corrected paths; full AG-3 was not rerun on the final source

pointwise Wilson 95% interval is used. Sensitivity-analysis output may also store a numerical Wilson interval for deterministic $N = 1$ cells, but such values are not interpreted as sampling uncertainty. Stored summaries from different sensitivity scenarios may be compared descriptively, but such comparisons are not treated as a new cross-scenario paired-effect estimator or post-hoc statistical test.

No additional hypothesis test, p -value, bootstrap procedure, or threshold defined after observing the results is introduced into the analysis. No interpolation is performed between processing-scale or contact-time evaluation points. Likewise, no post-hoc aggregate metric, family-wide mean, or maximum is constructed solely to explain selected results.

4.5 Validation and Evidence Scope

The processing implementation, measurement harness, production replay, and finite-contact projection address different error modes and are therefore not collapsed into a single validation result. Table 6 summarizes the property and implementation scope checked by each validation layer.

Processing-protocol correctness validation and the measurement harness’s untimed smoke serve different purposes. The former checks the normal protocol path, state, and final-key agreement of the study-specific Triple-KEM, whereas the latter confirms that the actual measurement harness completes the same

Table 7. Measured and Reconstructed Processing Times for the Seven Logical Stages

Stage	A' median (μs)	A' p05–p95 (μs)	B (μs)	$A' - B$ (μs)
G_SEND1	151.1570	150.5800–154.9111	150.6400	+0.5170
S_RECV1	88.5020	88.1290–92.0893	73.7200	+14.7820
S_SEND2	176.6835	176.1910–180.4661	176.6500	+0.0335
G_RECV2	132.3170	131.9420–135.9521	132.2320	+0.0850
G_SEND3_PRE	0.5350	0.5120–0.6070	0.7050	−0.1700
G_SEND3_POST	7.1690	7.1340–7.2110	7.1070	+0.0620
S_RECV3	7.5450	7.5100–7.5920	7.5910	−0.0460

protocol path normally and makes no timer calls in its smoke branch. Timed measurement applies `CLOCK_MONOTONIC_RAW` around the same stage calls, so functional validation results from correctness or smoke runs are not reclassified as official performance results.

The final CORE replay establishes same-source deterministic reproducibility for the SNAP-009 source and does not constitute independent implementation validation. The full direct comparison of the finite-contact projector was performed under SNAP-006; after two boundary-condition corrections, targeted direct-versus-projected comparisons and related regression and preflight checks were performed for the affected paths. No mismatch was observed in those targeted comparisons, but the full projector comparison was not rerun under final SNAP-009. Accordingly, the validation evidence is limited to the scope in Table 6 and does not establish independent implementation equivalence, equivalence to the original authors’ implementation, actual spacecraft processing performance, or independent external review.

5 Evaluation

5.1 Processing Characterization

The integrated stage measurement A' obtained in Section 4.1 serves as the stage-level reference used to construct the processing profile for T_1 . Across the seven stages, the medians range from 0.535 μs for G_SEND3_PRE to 176.6835 μs for S_SEND2, showing that the stage-level processing times are heterogeneous in the selected reference implementation. Table 7 reports the A' medians and empirical p05–p95 intervals together with the component reconstruction B and the signed residual from Eq. (12).

The values shown in Table 7 are obtained by converting the underlying ns values to μs and then rounding for display. The residuals are converted from the stored signed residuals rather than recomputed by subtracting the rounded table entries. The p05–p95 interval summarizes the empirical sample distribution of A' and is not a confidence interval for the median.

Figure 1(a) compares the A' medians and p05–p95 intervals in Table 7 with the reconstruction B . Because stage latencies span substantially different scales,




Fig. 1. Processing Characterization of the Seven Logical Stages

the latency axis is logarithmic. Figure 1(b) shows the signed residual $A' - B$ on a linear scale.

The residuals vary in both magnitude and sign across stages. The largest positive residual is $+14.782 \mu s$ for `S_RECV1`, whereas `G_SEND3_PRE` and `S_RECV3` have negative residuals of $-0.170 \mu s$ and $-0.046 \mu s$, respectively. Thus, summing separately measured component medians does not always reproduce the integrated stage measurement exactly. These residuals alone, however, do not identify a specific integration overhead, optimization effect, implementation element, or measurement error.

5.2 Canonical Measured-Processing Agreement

We first evaluate agreement between T_0 and T_1 under the canonical schedule, measured processing at $\alpha_S = 1$, $p^* = 2/5$, and `UNBOUNDED` contact. Each of the four stochastic CORE cells in Table 8 contains $N = 4096$ paired trials. No $R2$ trial is observed in any of the evaluated BCH or LDPC(512,256) TC-only or TM-only conditions, whereas $R1$ occurs in all four conditions.

For BCH, 311 and 51 $R1$ trials are observed in the TC-only and TM-only conditions, respectively; for LDPC(512,256), the corresponding counts are 43 and 13. Thus, under the evaluated canonical processing condition, equality of the four-class final transport outcome does not imply equality of the canonical event path. Conversely, observing $R2 = 0$ means only that no macro-outcome

Table 8. Canonical Agreement at $\alpha_S = 1$, $p^* = 2/5$, and UNBOUNDED Contact

TC code	Loss condition	N	$R0$	$R1$	$R2$
BCH	TC-only	4096	3785	311	0
BCH	TM-only	4096	4045	51	0
LDPC512	TC-only	4096	4053	43	0
LDPC512	TM-only	4096	4083	13	0

Table 9. Deterministic Timing Differences at the Canonical No-Loss $\alpha_S = 1$ Condition

Timing endpoint	$\Delta t_x = T_1 - T_0$ (μs)
$t_{\text{transport_complete}}$	+132.852
$t_{\text{ground_key_ready}}$	+140.021
$t_{\text{spacecraft_key_ready}}$	+140.397
$t_{\text{protocol_complete}}$	+140.397

discordance appears in the evaluated $N = 4096$ trials; it does not prove that the population discordance probability is exactly zero.

The three observation layers also remain distinct under deterministic no-loss conditions. Under the canonical schedule with $\alpha_S = 1$, the BCH and LDPC(512,256) no-loss controls are both $R0$, yet the continuous timing differences $T_1 - T_0$ are nonzero. The same values are observed for both TC channel codes, as shown in Table 9.

An $R0$ classification therefore does not mean that all observables are identical between the two models. It means that the four-class final outcome and canonical event path are equal, and it does not imply that both executions are successful. Because numerical timestamps are excluded from canonical-path equality, continuous timing differences such as those in Table 9 can remain even in $R0$ trials.

5.3 Processing-Scale Agreement

Using the eight code-specific α_S points defined in Section 4.3, we evaluate all 32 stochastic CORE cells under $p^* = 2/5$, the canonical schedule, and UNBOUNDED contact. Figure 2 shows the $R0/R1/R2$ composition over the full predefined processing-scale grid for BCH and LDPC(512,256) under TC-only and TM-only loss. Each stochastic cell contains $N = 4096$ paired trials. Each stacked bar corresponds to one predefined α_S evaluation point, and the categorical horizontal spacing does not represent numerical distance between α_S values. No interpolation is performed between evaluation points.

No $R2$ trial is observed in any stochastic TC-only cell. For BCH, $R1$ occurs at every evaluated point, and the $R0/R1$ composition changes with processing scale. For LDPC(512,256), all 4096 pairs are $R0$ at $\alpha_S = 0$, while $R1$ occurs at every subsequently evaluated processing-scale point. Thus, over the evaluated




Fig. 2. Processing-Scale Agreement over the Predefined Grid

TC-only grid, canonical-path agreement composition changes with processing scale, while four-class macro-outcome discordance is not observed.

Under TM-only loss, cells with $R2 = 0$ are observed at lower processing-scale points, while compositions containing $R2$ appear at the code-specific landmarks and at some larger predefined points. At L_{BCH} , BCH has $R0/R1/R2 = 0/3854/242$; at L_{LDPC} , LDPC(512,256) has $0/3830/266$. These observations apply only to the predefined discrete grid and do not define L_{BCH} or L_{LDPC} as continuous processing-time thresholds.

Within a given code, the TC-only and TM-only families share the same eight-point α_S grid, so conditions at the same processing scale can be compared descriptively. At L_{BCH} , BCH has 2356/1740/0 in TC-only versus 0/3854/242 in TM-only. At L_{LDPC} , LDPC(512,256) has 1544/2552/0 in TC-only versus 0/3830/266 in TM-only. These are descriptive comparisons between evaluated

Table 10. Bidirectional Macro-Outcome Transitions at the TM-Only Landmark Cells

TC code	α_S	FP2→ SUCCESS	SUCCESS→ FP2
BCH	L_{BCH}	124	118
LDPC512	L_{LDPC}	138	128

Table 11. Finite-Contact Agreement for the Selected TM-Only Landmark Scenarios

Scenario	Predefined evaluation point(s)	$R0$	$R1$	$R2$
BCH TM-only, $\alpha_S = L_{\text{BCH}}$	$H = 1.0$ s	4096	0	0
	H indices 2–16 (each predefined point from 1.8 to 500.0 s)	0	3854	242
	UNBOUNDED	0	3854	242
LDPC512 TM-only, $\alpha_S = L_{\text{LDPC}}$	Each of $H = 1.0, 1.8, 2.4$ s	4096	0	0
	H indices 4–16 (each predefined point from 3.1 to 500.0 s)	0	3830	266
	UNBOUNDED	0	3830	266

conditions that share the same code and processing scale but differ in the loss direction; they do not establish a universal TC/TM sensitivity ordering or a causal effect of a single transport mechanism.

Because the $R2$ composition alone does not show the direction of macro-outcome changes, Table 10 reports the four-class transitions in the two TM-only landmark cells that contain $R2$. Both cells use TM-only loss, $p^* = 2/5$, the canonical schedule, UNBOUNDED contact, and $N = 4096$; transition directions are $T_0 \rightarrow T_1$.

The only nonzero off-diagonal four-class transitions in the two landmark cells are the two directions between FP2 and SUCCESS. The evidence therefore does not support an interpretation in which adding processing time either preserves success status for every paired trial or changes it only toward failure. These counts describe pair-level transition directions and are not an analysis of aggregate success-rate effects or statistical significance.

5.4 Finite-Contact Agreement

Finite-contact projection is applied to all 48 CORE scenarios at the 16 horizons defined in Section 4.3. The main text reports in detail the two TM-only landmark scenarios that connect directly to the processing-scale results. The BCH condition uses $\alpha_S = L_{\text{BCH}}$, the LDPC(512,256) condition uses $\alpha_S = L_{\text{LDPC}}$, and both use $p^* = 2/5$. Each evaluation point contains $N = 4096$ paired trials, and counts in a grouped row of Table 11 apply separately to every horizon listed in that row.

For BCH, all pairs are $R0$ at $H = 1.0$ s, whereas each predefined finite horizon from 1.8 through 500.0 s contains 3854 $R1$ and 242 $R2$ pairs. For LDPC(512,256), all pairs are $R0$ at each of 1.0, 1.8, and 2.4 s, whereas every predefined finite horizon from 3.1 s onward contains 3830 $R1$ and 266 $R2$ pairs. Here, $R0$ means that the two models have the same four-class outcome and canonical path; it does not mean that the trial is successful.

These observations are limited to the predefined discrete contact-time grid. Accordingly, neither 1.8 s nor 3.1 s is interpreted as a continuous contact-time threshold. The 16 horizons for a given scenario are projections of the same underlying paired executions and are therefore not treated as independent Monte Carlo samples. **UNBOUNDED** is a separate reference from 500.0 s.

The supporting classification assigns the finite $R1/R2$ groups under these conditions to categories that include canonical-path divergence. The corresponding finite $R2$ groups are also classified as having macro-outcome discordance in the contact-unbounded execution. These classifications delimit the interpretation of the finite-contact results and do not show that a particular contact horizon newly creates the discordance or identify its causal mechanism.

5.5 Schedule and Model-Assumption Sensitivities

TM-Phase Existence Case The TM-phase sensitivity contains an existence case in which a processing-time difference appears together with a larger timing difference aligned with the schedule structure. The stochastic case reported here uses LDPC(512,256), TM-only loss, $p^* = 2/5$, $\alpha_S = 1$, canonical CLCW phase $k = 0$, and TM phase $k = 5$. All $N = 4096$ paired trials in this cell have different canonical event paths. For the binary outcome, 144 trials change from success in T_0 to failure in T_1 , 115 change in the reverse direction, 197 are successful in both models, and 3640 fail in both models.

Separately from the stochastic path/binary result, we compare the phase-specific model timing difference in a no-loss deterministic control with the same LDPC(512,256) code and $\alpha_S = 1$. At canonical TM phase $k = 0$,

$$\Delta t_{\text{transport_complete}} = 0.132852 \text{ ms}$$

whereas at TM phase $k = 5$,

$$\Delta t_{\text{transport_complete}} = 5.538852 \text{ ms.}$$

Each value is the transport-completion timing difference between T_1 and T_0 at the corresponding TM phase. Their difference,

$$5.538852 - 0.132852 = 5.406 \text{ ms,}$$

is the phase-dependent change in the model-to-model timing difference; it does not mean that the absolute transport-completion time of T_1 increases by 5.406 ms when the phase changes. This value is numerically equal to the reference TM slot duration T_{slot} .

This numerical correspondence is consistent with an example of schedule-aligned timing amplification, but it does not prove that one particular TM slot was causally missed. The $k = 5$ case was selected after analysis as an existence case and is not interpreted as representative, typical, worst-case, or a phase threshold.

CLCW-Phase Scope Check The CLCW-phase sensitivity varies only the CLCW phase over eight evaluation points while fixing the TM phase at canonical $k = 0$. Both directions of binary discordance are zero in all 48 evaluated CLCW-phase scenarios. However, under some stochastic code/loss conditions, the canonical-event-path disagreement count varies with CLCW phase. Thus, preservation of binary success/failure agreement over the evaluated CLCW-phase conditions does not imply that the canonical path is phase independent.

For each code/loss combination, the stored support, median, q05, and q95 summaries for all four timing endpoints are the same at every CLCW phase as in the corresponding $k = 0$ summary. Identical timing summaries, however, do not prove that the complete paired Δt sample sequences are identical. Because this sensitivity fixes the TM phase at $k = 0$, it does not evaluate whether the preceding TM $k = 5$ existence case persists under other CLCW phases.

P1 Readiness at Contact Open The P1-ready-at-open sensitivity assumes, unlike the baseline STRICT_ONLINE condition, that a fresh P1 and matching ground state are already prepared at contact open. Both directions of binary discordance are zero in all six scenarios of this sensitivity. The canonical-path disagreement count nevertheless differs from the baseline CORE condition in some stochastic cases.

For BCH TC-only, the count changes from 311 in the canonical $\alpha_S = 1$ CORE condition to 185 under ready-at-open, and for BCH TM-only from 51 to 13. For LDPC(512,256) TC-only, the count changes from 43 to 14, while both LDPC(512,256) TM-only conditions contain 13 path disagreements. No path disagreement is observed in either TC code under the corresponding no-loss baseline or ready-at-open execution. The stored support and median/q05/q95 summaries for the four timing endpoints are the same as in the corresponding canonical CORE scenarios.

The canonical processing result is therefore not interpreted as a universal path property independent of whether P1-generation processing is already complete at contact open. At the same time, equality of the timing summaries between the ready-at-open and canonical CORE conditions does not establish equality of their absolute event times, pair-level traces, or complete timing distributions.

Empirical Processing Resampling The processing-resampling sensitivity replaces the fixed-median processing profile with per-trial resampling of the measured empirical ground and spacecraft endpoint vectors. Table 12 compares the canonical-path disagreement counts for the fixed-median CORE and resampling sensitivity at $\alpha_S = 1$ under the canonical schedule.

Table 12. Canonical-Path Disagreement under Fixed-Median and Empirical Processing Inputs

TC code	Loss condition	Fixed-median path unequal	Resampled path unequal
BCH	no-loss	0/1	1/4096
BCH	TC-only	311/4096	314/4096
BCH	TM-only	51/4096	52/4096
LDPC512	no-loss	0/1	0/4096
LDPC512	TC-only	43/4096	43/4096
LDPC512	TM-only	13/4096	13/4096

Both directions of binary discordance are zero in all six resampling scenarios. The stored timing summaries for the resampling scenarios, however, are not identical to those of their corresponding fixed-median scenarios. Thus, even with empirical processing inputs, binary success/failure agreement does not imply equality of the canonical path or timing.

The no-loss rows in Table 12 require separate interpretation. The fixed-median no-loss cell is one deterministic control, whereas the resampling no-loss cell is a stochastic sensitivity with $N = 4096$ because the processing vector itself is resampled. Their raw counts are therefore not compared as probability estimates from identical sampling designs. This sensitivity characterizes the empirical variability in the measured endpoint vectors and is not a general robustness evaluation for arbitrary processor variability or actual spacecraft hardware.

5.6 Discussion

The evaluation shows that omitting processing time can affect final transport outcome, canonical event path, and continuous timing differently. Under the canonical stochastic condition at $\alpha_S = 1$, $R1$ exists even when no $R2$ trial is observed, and continuous timing differences remain in deterministic no-loss $R0$ cases. The three observation layers therefore provide different execution information rather than an ordered scale of severity, and $R0$, $R1$, and $R2$ should be interpreted as paired classifications constructed from Q and Π .

The adequacy of a zero-processing abstraction can therefore depend on the objective of the analysis. For an outcome-only analysis, some $R1$ differences may leave the required observable unchanged. If retransmission, CLCW feedback, timer behavior, or schedule interaction is itself part of the analysis, canonical-path equality can be relevant. Likewise, when subsequent analysis depends on continuous times such as key-ready or transport-completion time, residual Δt_x in $R0$ cases should not be assumed negligible by definition.

The processing-scale evaluation shows the same distinction. No $R2$ trial is observed in any evaluated TC-only stochastic cell, while $R1$ is observed at every point except LDPC(512,256) at $\alpha_S = 0$. The TM-only grid contains both compositions with $R2 = 0$ and compositions containing $R2$ at different predefined

processing-scale points. TC-only and TM-only conditions can be compared at the same code and α_S , but the observed differences are conditional on the transport instantiation, $p^* = 2/5$, the canonical schedule, and the predefined processing grid used here. They are not generalized as a universal sensitivity ordering between physical TC and TM channels or as evidence for a single causal mechanism.

Finite-contact results show that the observed agreement composition can also depend on the evaluated contact horizon under the same processing configuration. The schedule and model-assumption sensitivities likewise show that the agreement state at different observation layers can vary with TM phase, CLCW phase, P1 readiness, and empirical processing input. These results are restricted to the predefined discrete conditions and executed scope and do not extend to unevaluated phase interactions or continuous parameter regions.

An observed canonical-path difference means that the modeled transport and control event history differs between the two executions. Path disagreement itself does not measure CPU or energy consumption, communication-resource cost, operational damage, or security failure. Resource or security consequences that are not directly observed in this study are therefore not inferred from *R1* or *R2*.

5.7 Limitations

The processing measurements are limited to one Intel x86-64 reference platform and the OpenSSL-based Triple-KEM implementation developed for this study. ML-KEM-512 is used as a processing approximation for the Kyber512-based transport input, and the measured implementation is not the original implementation used by Hülsing et al. The 5,000 samples are complete-handshake samples collected in one measurement run rather than repetitions across different machines, dates, or independent measurement campaigns. The measured processing profile is therefore not generalized directly to flight hardware or other processor architectures.

The parameter α_S is an analytical scaling factor for the measured spacecraft-side processing vector rather than an actual spacecraft-processor CPU-speed ratio. The code-specific landmarks are evaluation reference points obtained by the deterministic search procedure, not boundaries between processor classes. The production evaluation uses only eight predefined discrete α_S points, so no continuous threshold or monotonic relationship is inferred between adjacent points.

The loss model abstracts channel errors as independent Bernoulli losses of complete post-decoding TC CLTUs or TM Transfer Frames. The value $p^* = 2/5$ is a stress condition selected for the stochastic evaluation by the predefined pilot-selection rule and does not represent the loss statistics of an actual mission channel. Temporally correlated errors, burst loss, physical-layer bit-error processes, and mission-specific link adaptation are not evaluated.

The main stochastic evaluation uses one reference TM transport profile and two TC channel codes. Finite-contact analysis uses only the 16 predefined horizons, and the horizons of a given scenario are projections of the same underlying paired

executions rather than independent Monte Carlo samples. UNBOUNDED is a separate reference from 500.0 s.

TM phase and CLCW phase are each evaluated one factor at a time, and the 8×8 interaction between the two axes is not executed. Results from one phase sensitivity therefore cannot be generalized to combinations with the other phase axis. The TM $k = 5$ example in the main text is a post-hoc existence case and was not selected as a representative, typical, or worst-case phase.

The processing-resampling sensitivity draws the ground four-stage vector and spacecraft three-stage vector separately. It therefore preserves joint sample structure within each endpoint but not the full ground-spacecraft dependence present in an original complete-handshake row. This sensitivity is a supporting evaluation of empirical variability in the collected reference measurements, not a model of arbitrary hardware-processing variability.

The P1-ready-at-open sensitivity evaluates only the specific readiness assumption that a fresh P1 and matching ground state are available at contact open. It does not remove the G_SEND1 computation itself, set all ground processing to zero, or evaluate long-term precomputation or P1 reuse across sessions.

The component reconstruction B has different measurement boundaries and execution context from the integrated measurement A' . Consequently, the $A' - B$ residual does not provide an overhead decomposition, independent validation, or causal estimate of a specific optimization effect. Likewise, the relationship between B and the percentile interval of A' is not interpreted as a statistical-significance or equivalence test.

The validation evidence is also scope limited. The final CORE replay establishes deterministic reproducibility from the same source but not equivalence to an independent implementation. The full direct comparison of the finite-contact projector was performed under SNAP-006; after two boundary-condition corrections, targeted direct-versus-projected regression and related preflight checks were performed, but the same full projector comparison was not repeated under final SNAP-009. The overall validation campaign also did not undergo independent external review.

Within these limitations, the evaluation supports the conclusion that, for the evaluated PQC-based OTAR transport model, the agreement state under processing-time omission depends on whether the observation level is final transport outcome, canonical event path, or continuous timing, and the result must be described together with the corresponding processing scale, loss direction, contact-time, and schedule conditions.

6 Conclusion

This study evaluated the effect of omitting processing time when modeling PQC-based OTAR over CCSDS TC/TM links by comparing a zero-processing baseline T_0 with a processing-aware model T_1 . The comparison separates three observation levels: final transport outcome, canonical event path, and continuous timing. The P1-P2-P3 exchange of a PQC-based SDLS key update is instantiated together

with TC/TM transport, COP-1 reliability control, CLCW return feedback, finite contact time, and direction-specific loss, while T_1 incorporates processing times measured at seven logical stages of a study-specific Triple-KEM reference implementation.

Under the canonical schedule with measured processing at $\alpha_S = 1$, no final-transport-outcome discordance is observed in the evaluated BCH and LDPC(512,256) TC-only or TM-only conditions, although some paired trials have different canonical event paths. Continuous timing differences also remain in deterministic $R0$ cases, so equality of the final transport outcome and event path does not imply equality of timing information. At some predefined processing-scale points, additional $R2$ trials appear under TM-only loss, with transitions occurring in both directions between FP2 and SUCCESS. By contrast, selected higher-scale TC-only conditions contain many event-path disagreements while still showing no $R2$ trial.

Finite-contact projection shows that the agreement composition can vary with the evaluated contact time even when the processing configuration is unchanged. The schedule-phase analysis also identifies an existence case in which the processing-time difference appears together with a larger transport-completion timing difference numerically aligned with the TM slot duration. These observations are limited to the predefined discrete contact-time and schedule conditions and do not establish continuous thresholds, monotonicity, or a causal missed-slot mechanism.

Overall, under the evaluated conditions, the effect of omitting processing time cannot be characterized adequately by whether the key-update exchange succeeds or fails. Instead, the relevant agreement level—final transport outcome, canonical event path, or continuous timing—must be stated together with the processing scale, contact time, loss direction, and schedule conditions under which the result is observed. The numerical results are conditional on the selected transport configuration, processing implementation, x86-64 reference measurement environment, and predefined evaluation grid and should not be generalized directly to operational spacecraft hardware or unevaluated operating environments.

References

1. Consultative Committee for Space Data Systems: Communications operation procedure-1. Tech. Rep. CCSDS 232.1-B-2, CCSDS, Washington, DC (Sep 2010), including Technical Corrigendum 1, April 2019
2. Consultative Committee for Space Data Systems: Space data link security protocol—extended procedures. Tech. Rep. CCSDS 355.1-B-1, CCSDS, Washington, DC (Feb 2020)
3. Consultative Committee for Space Data Systems: Tc space data link protocol. Tech. Rep. CCSDS 232.0-B-4, CCSDS, Washington, DC (Oct 2021), including Technical Corrigendum 1, October 2023
4. Consultative Committee for Space Data Systems: Tm space data link protocol. Tech. Rep. CCSDS 132.0-B-3, CCSDS, Washington, DC (Oct 2021)
5. Consultative Committee for Space Data Systems: Space data link security protocol. Tech. Rep. CCSDS 355.0-B-2, CCSDS, Washington, DC (Aug 2022)

6. Ghosh, T., Nath, I.: Secure satellite communication in the post-quantum era: A lattice-based cryptographic approach. *International Journal of Satellite Communications and Networking* **44**(5), 524–543 (2026). <https://doi.org/10.1002/sat.70041>
7. Hülsing, A., Lange, T., Weber, F.J.: A key-update mechanism for the space data link security protocol. In: *Cryptology and Network Security. Lecture Notes in Computer Science*, vol. 16351, pp. 602–611. Springer (2026). https://doi.org/10.1007/978-981-95-4434-9_29
8. Kim, D., Jang, J.: Satellite OTAR technology classification and directions for future research. *Journal of the Korea Institute of Information Security and Cryptology* **36**(1), 13–23 (Feb 2026). <https://doi.org/10.13089/JKIISC.2026.36.1.13>
9. Mutlugun, A., Hanna, Y., Akkaya, K.: Performance evaluation of quantum-resistant IKEv2 protocol for satellite networking environments. In: *2024 IEEE Virtual Conference on Communications (VCC)*. pp. 1–7. IEEE (2024). <https://doi.org/10.1109/VCC63113.2024.10914395>
10. Wilson, E.B.: Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association* **22**(158), 209–212 (1927). <https://doi.org/10.1080/01621459.1927.10502953>